

What is claimed is:

1 1. A method of selectively caching content responsive to a cache miss, comprising steps of:
2 receiving, at a cache store responsive to a cache miss, content for which the cache miss
3 occurred;
4 deciding whether the received content should be cached at the cache store, responsive to
5 the receiving step, and only caching it if so; and
6 returning the received content from the cache store to a client that sent a request that
7 caused the cache miss, regardless of the deciding step.

1 2. The method according to Claim 1, wherein the deciding step evaluates historical metrics.

1 3. The method according to Claim 1, wherein the deciding step further comprises evaluating
2 a hit rate associated with the content and deciding whether content having that hit rate may be
3 advantageously cached by the cache store.

1 4. The method according to Claim 1, wherein the deciding step further comprises deciding
2 whether a hit rate associated with the content is higher than hit rates associated with other content
3 already cached by the cache store and if so, deciding to accept the content.

1 5. The method according to Claim 1, wherein the deciding step considers historical metrics
2 associated with the content.

- 1 6. The method according to Claim 1, wherein the deciding step considers resources of the
2 cache store.
- 1 7. The method according to Claim 1, wherein the deciding step considers currently-available
2 resources of the cache store.
- 1 8. The method according to Claim 1, wherein the deciding step compares a priority
2 associated with the content to priorities associated with already-cached content at the cache store.
- 1 9. A system for selectively caching content responsive to a cache miss, comprising:
2 means for receiving, at a cache store responsive to a cache miss, content for which the
3 cache miss occurred;
4 means for deciding whether the received content should be cached at the cache store,
5 responsive to the means for receiving, and only caching it if so; and
6 means for returning the received content from the cache store to a client that sent a
7 request that caused the cache miss, regardless of an outcome of the means for deciding.
- 1 10. The system according to Claim 9, wherein the means for deciding further comprises means
2 for evaluating a hit rate associated with the content and deciding whether content having that hit
3 rate may be advantageously cached by the cache store.
- 1 11. The system according to Claim 9, wherein the means for deciding further comprises means

for deciding whether a hit rate associated with the content is higher than hit rates associated with other content already cached by the cache store and if so, deciding to accept the content.

12. The system according to Claim 9, wherein the means for deciding considers one or more of: historical metrics associated with the content; resources of the cache store; and currently-available resources of the cache store.

13. The system according to Claim 9, wherein the means for deciding compares a priority associated with the content to priorities associated with already-cached content at the cache store.

14. A computer program product for selectively caching content responsive to a cache miss, the computer program product embodied on one or more computer-readable media and comprising:

computer-readable program code means for receiving, at a cache store responsive to a cache miss, content for which the cache miss occurred;

computer-readable program code means for deciding whether the received content should be cached at the cache store, responsive to the computer-readable program code means for receiving, and only caching it if so; and

computer-readable program code means for returning the received content from the cache store to a client that sent a request that caused the cache miss, regardless of an outcome of the computer-readable program code means for deciding.

1 15. The computer program product according to Claim 14, wherein the computer-readable
2 program code means for deciding further comprises computer-readable program code means for
3 evaluating a hit rate associated with the content and deciding whether content having that hit rate
4 may be advantageously cached by the cache store.

1 16. The computer program product according to Claim 14, wherein the computer-readable
2 program code means for deciding further comprises computer-readable program code means for
3 deciding whether a hit rate associated with the content is higher than hit rates associated with
4 other content already cached by the cache store and if so, deciding to accept the content.

1 17. The computer program product according to Claim 14, wherein the computer-readable
2 program code means for deciding considers one or more of: historical metrics associated with the
3 content; resources of the cache store; and currently-available resources of the cache store.

1 18. The computer program product according to Claim 14, wherein the computer-readable
2 program code means for deciding compares a priority associated with the content to priorities
3 associated with already-cached content at the cache store.